

(19) World Intellectual Property Organization  
International Bureau(43) International Publication Date  
2 August 2001 (02.08.2001)


PCT

(10) International Publication Number  
**WO 01/055856 A3**

- (51) International Patent Classification<sup>7</sup>: G06F 3/06, 17/30 (74) Agent: PARKER, Lee; Commvault Systems, Inc., 2 Crescent Place, Suite B, Oceanport, NJ 07757-0900 (US).
- (21) International Application Number: PCT/US01/03183 (81) Designated State (*national*): CA.
- (22) International Filing Date: 31 January 2001 (31.01.2001) (84) Designated States (*regional*): European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR).
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
60/179,344 31 January 2000 (31.01.2000) US  
09/774,268 30 January 2001 (30.01.2001) US
- (71) Applicant: COMMVAULT SYSTEMS, INC. [US/US];  
2 Crescent Place, Oceanport, NJ 07757-0900 (US).
- (72) Inventors: IGNATIUS, Paul; 25 Highland Drive, Jackson, NJ 08527 (US). THEISEN, Marjorie, H.; 3272 S. Newcombe Street, Lakewood, CO 80227 (US). OSHINSKY, David, A.; 22 Francis Road, East Brunswick, NJ 08816 (US). KAVURI, Srinivas; 40 Maple Court, Highland Park, NJ 08904 (US).
- Published:**  
 — with international search report  
 — before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments
- (88) Date of publication of the international search report:  
4 December 2003
- (15) Information about Correction:  
Previous Correction:  
see PCT Gazette No. 44/2002 of 31 October 2002, Section II

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: LOGICAL VIEW AND ACCESS TO PHYSICAL STORAGE IN MODULAR DATA AND STORAGE MANAGEMENT SYSTEM



(57) Abstract: A computer storage system having a processor that supports operation of at least one software application in order to store selected data in the computer storage system, at least one storage media for storing the selected data, a plurality of storage policies, each having particular storage guidelines, that are available to determine how data is to be stored in the computer storage system, and a storage control that interacts with the at least one software application of the processor to determine which of the plurality of storage policies to use for storage of the selected data and that stores the selected data according to the selected storage policy. In one embodiment, the storage control of the computer storage system migrates the selected data from one of the at least one storage media to another of the at least one storage media according to the selected storage policy.

**WO 01/055856 A3**

## INTERNATIONAL SEARCH REPORT

Int'l Application No  
PCT/US 01/03183

A. CLASSIFICATION OF SUBJECT MATTER  
IPC 7 G06F3/06 G06F17/30

According to International Patent Classification (IPC) or to both national classification and IPC

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, PAJ, INSPEC

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 0 670 543 A (IBM) 6 September 1995 (1995-09-06) column 7, line 9-40 column 8, line 41 -column 9, line 49 claims 1,7,8 figures 2,6A	1-4,9-12
A	EP 0 341 230 A (IBM) 8 November 1989 (1989-11-08) column 1, line 32 - line 53 column 6, line 25 - line 53 claim 1	5-8 1-12

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

## \* Special categories of cited documents:

- \*A\* document defining the general state of the art which is not considered to be of particular relevance
- \*E\* earlier document but published on or after the international filing date
- \*L\* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- \*O\* document referring to an oral disclosure, use, exhibition or other means
- \*P\* document published prior to the international filing date but later than the priority date claimed

- \*T\* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- \*X\* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- \*Y\* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- \*&\* document member of the same patent family

Date of the actual completion of the international search

25 September 2003

Date of mailing of the international search report

02/10/2003

Name and mailing address of the ISA

European Patent Office, P.B. 5618 Patentlaan 2  
NL - 2280 HV Rijswijk  
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,  
Fax: (+31-70) 340-3016

Authorized officer

Perrier, S

## INTERNATIONAL SEARCH REPORT

Information on patent family members

In - al Application No

PCT/US 01/03183

Patent document cited in search report		Publication date		Patent family member(s)	Publication date
EP 0670543	A	06-09-1995	US	5491810 A	13-02-1996
			EP	0670543 A1	06-09-1995
			JP	7281940 A	27-10-1995
			US	5790886 A	04-08-1998
EP 0341230	A	08-11-1989	EP	0341230 A2	08-11-1989
			JP	1292452 A	24-11-1989



## CORRECTED VERSION

(19) World Intellectual Property Organization  
International Bureau(43) International Publication Date  
2 August 2001 (02.08.2001)(10) International Publication Number  
**WO 01/055856 A2**

PCT


- (51) International Patent Classification<sup>7</sup>: **G06F 12/00**
- (21) International Application Number: PCT/US01/03183
- (22) International Filing Date: 31 January 2001 (31.01.2001)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
60/179,344 31 January 2000 (31.01.2000) US  
09/774,268 30 January 2001 (30.01.2001) US
- (71) Applicant: COMMVAULT SYSTEMS, INC. [US/US];  
2 Crescent Place, Oceanport, NJ 07757-0090 (US).
- (72) Inventors: IGNATIUS, Paul; 25 Highland Drive, Jackson, NJ 08527 (US). THEISEN, Marjorie, H.; 3272 S. Newcombe Street, Lakewood, CO 80227 (US). OSHINSKY, David, A.; 22 Francis Road, East Brunswick, NJ 08816 (US). KAVURI, Srinivas; 40 Maple Court, Highland Park, NJ 08904 (US).
- (74) Agent: PARKER, Lee; Commvault Systems, Inc., 2 Crescent Place, Suite B, Oceanport, NJ 07757-0900 (US).
- (81) Designated State (*national*): CA.
- (84) Designated States (*regional*): European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR).
- Published:**  
— without international search report and to be republished upon receipt of that report
- (48) Date of publication of this corrected version:  
31 October 2002

## (15) Information about Correction:

see PCT Gazette No. 44/2002 of 31 October 2002, Section II

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: LOGICAL VIEW AND ACCESS TO PHYSICAL STORAGE IN MODULAR DATA AND STORAGE MANAGEMENT SYSTEM



WO 01/055856 A2

(57) Abstract: A computer storage system having a processor that supports operation of at least one software application in order to store selected data in the computer storage system, at least one storage media for storing the selected data, a plurality of storage policies, each having particular storage guidelines, that are available to determine how data is to be stored in the computer storage system, and a storage control that interacts with the at least one software application of the processor to determine which of the plurality of storage policies to use for storage of the selected data and that stores the selected data according to the selected storage policy. In one embodiment, the storage control of the computer storage system migrates the selected data from one of the at least one storage media to another of the at least one storage media according to the selected storage policy.

**LOGICAL VIEW AND ACCESS TO PHYSICAL STORAGE IN  
MODULAR DATA AND STORAGE MANAGEMENT SYSTEM**

**I. SPECIFICATION**

**CROSS-REFERENCE TO RELATED APPLICATION**

The present application claims the benefit of U.S. Provisional Application Serial Number 60/179,344, filed January 31, 2000, pending, which is hereby incorporated by reference herein in its entirety.

**BACKGROUND**

**1. Field of the Invention**

The present invention generally relates to data storage in computer systems and more particularly to the logical view and access to physical storage in a modular data and storage management system.

**2. Description of the Related Art**

Existing storage systems require an end user to know and define exactly where data needs to be written and deposited in the storage system. This type of storage system is very hard to use and inherently unscalable as the user's data grows since the user will not be able to keep track of every storage device in the system. More importantly, every time new storage device(s) are added to the environment, the user will have to redefine how and where archival or backup data needs to be stored.

Many other problems and disadvantages of the prior art will become apparent to one skilled in the art after comparing such prior art with the present invention as described herein.

**SUMMARY**

Various aspects of the present invention may be realized through a computer storage system that includes a processor that supports operation of at least one software application that may be utilized to generate data that is stored in the computer storage system. The computer storage system also includes at least one storage media for storing selected data from the data that is generated. A plurality of storage policies are included where each one of the plurality of storage policies includes directions that the processor must follow for storage of data on the at least one storage media. A storage control interacts with the at least one software application of the processor to select one of the

plurality of storage policies and to determine, according to the selected one of the plurality of storage policies, which of the at least one storage media to use for storage of the selected data.

In certain embodiments, the at least one storage media of the computer storage system includes a plurality of storage media, e.g., magnetic tape media, magnetic disk media, optical medial, or other suitable storage media familiar to those skilled in the art. The storage control of the computer storage system may include a manager module that interacts with a media module to compile an index to track the location of the selected data in the at least one storage media. The storage control is often configured to migrate the selected data among the at least one storage media according to the selected one of the plurality of storage policies. The plurality of storage policies may be storage sequences, storage logic, initial storage sequence selection criteria, storage sequence reselection criteria, storage sequence adaptation criteria, etc. The initial storage sequence selection criteria are commonly user directed override, user profile, application, file type, user network location, and available storage space or similar criteria. The storage sequence reselection criteria are criteria such as specific file usage history, file type usage history, user profile, user network relocation, available storage space, added storage media, etc. Finally, the storage sequence adaptation criteria are items such as specific file usage history, user profile, user network relocation, available storage space, added storage media, or other suitable criteria that becomes apparent to those of ordinary skill in the art and viewing the present disclosure.

Various aspects of the present invention may also be realized by a computer storage system having a processor that supports operation of at least one software application in order to store selected data in the computer storage system, at least one storage media for storing the selected data, a plurality of storage policies, each having particular storage guidelines, that are available to determine how data is to be stored in the computer storage system, and a storage control that interacts with the at least one software application of the processor to determine which of the plurality of storage policies to use for storage of the selected data and that stores the selected data according to the selected storage policy. In one embodiment, the storage control of the computer storage system migrates the selected data from one of the at least one storage media to another of the at least one storage media according to the selected storage policy.

Still other aspects of the present invention are realized through a method for storing data on one of a plurality of storage media of a computer storage system according to one of a plurality of storage policies that are defined by a user. The method involves, not necessarily in this order directing a software application to store selected data; examining the selected data for particular characteristics; selecting a particular storage policy that matches the particular characteristics of the

selected data; and storing the selected data on a storage media of the computer storage system according to the particular storage policy.

The method may also include migrating the selected data among the plurality storage media according to the particular storage policy.

Other aspects of the present invention will become apparent with further reference to the drawings and specification which follow.

### BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram of an exemplary embodiment of a storage system that operates according to principles of the present invention.

Fig. 2 is a block diagram of an exemplary storage system illustrated in greater detail than the storage system of Fig. 1.

Fig. 3 is a block diagram of another exemplary storage system that is illustrated in greater detail than the storage system of Fig. 1.

Fig. 4 is a block diagram of another exemplary storage system that illustrates another embodiment of the storage system of Fig. 1.

Fig. 5 is a block diagram of exemplary storage policies as illustrated in Fig. 1.

Fig. 6 is a block diagram of exemplary storage sequences in greater detail than as illustrated in Fig. 5.

Fig. 7 is a block diagram of the initial storage sequence selection illustrated in greater detail than in Fig. 5.

Fig. 8 is a block diagram of the storage sequence reselection illustrated in greater detail than in Fig. 5.

Fig. 9 is a block diagram of the storage sequence adaptation illustrated in greater detail than in Fig. 5.

### DETAILED DESCRIPTION OF THE DRAWINGS

An exemplary description of the principles according to the present invention may be found in a storage policy where logical data buckets are presented to the end user for directing their data storage operations such as backup and archival. Storage policy is a logical concept. Each storage policy comprises one or more copies. Each copy is a self-contained unit of information. A primary copy indicates the default destination of storage operations and is tied to a practical set of drives, e.g., physical drives. Drives are addressed independently of the library or media agent that they are attached to. Storage policy at the same time comprises of a multiple number of streams which indicate the number of active drives that the storage operation can run against, allowing parallelism in data transfer. A single copy backup may always be run which transfers data to the primary copy. The primary copy may then be replicated to the secondary copies or a multiple copy simultaneous backup may be performed.

Each copy could have its destination media type different to allow storage policies to hide media type and related problems in a clean abstract interface. It also takes care of keeping track of which copy resides where, what is size of each data, etc. In this manner, the user's data is protected as well as presented in an abstract way using storage policies. Each of the copies of storage policy could be governed with different aging rules for the application's data. For example, the primary copy could be targeted to magnetic disc and aged every 30 days and secondary copy could be targeted to tape and aged every year. Data read operation will always fetch it from the optimal solution for retrieval.

In one embodiment of the present invention, the user just has to define a "Storage Policy" which defines where, how, and the duration the data should be stored at a higher level of abstraction without having to have intimate knowledge of the underlying storage and technology. The management of the details of data storage is transparent to the user.

Fig. 1 is a block diagram of an exemplary embodiment of a storage system 100 that operates according to principles of the present invention. The storage system 100 includes software applications 102 that a user interacts with to store data to various storage media. The software applications 102 communicate with a storage control 104 that determines where to store data from the software applications 102 according to storage policies 106. The storage control 104 typically receives data that is to be stored from the software applications 102. Upon receipt of data from the software applications 102, the storage control 104 determines whether to store the data in a first storage media 108, a second storage media 110, or any number of other storage media down to an Nth

storage media 112. The storage control 104 controls which storage media the data will be stored to as well as where data that has been previously stored will be moved to.

For example, the storage control 104 may move data that has been stored in the first storage media 108 into the second storage media 110 based on certain storage policies 106. In addition, the storage control 104 could move data from the second storage media 110 to any other of the storage media including the Nth storage media 112. Also, the storage control 104 could move data from the first storage media 108 into the Nth storage media 112. Of course, the data could be moved in either direction, i.e., the storage control 104 is capable of moving data between storage media.

Fig. 2 is a block diagram of an exemplary storage system 200 illustrated in greater detail than the storage system 100. The storage system 200 includes a single computing device 202 to carry out operations of the storage system 200. The computing device 202 includes a software application 204 that interacts with an installed file system 206 when a user desires to transfer data from the software application 204 to the storage media. For example, when a user saves data from the software application 204, the user is only aware that the data is being saved to the installed file system 206. The installed file system 206 interacts with a manager module 208, which in turn interacts with a media module 210.

The manager module 208 includes storage policies 212. The storage policies 212 determine how data that the manager module 208 received from the installed file system 206 will be stored in the storage media. The storage policies 212 interact with creation tools 214 to create the appropriate storage scenario for the data that the manager module 208 has received from the installed file system 206. The creation tools 214 may be accessed directly in the manager module 208, or remotely, as illustrated by remote creation tools 216. The manager module 208 also includes a master map 218 to assist in initial storage decisions in the storage media.

The media module 210 includes a data index 220 that includes further details of where the data is to be stored in the storage system 200. The data index 220 includes details such as the location of magnetic disc media 222 and magnetic tape media 224. The data index 220 is updated with file location information when any data is moved from one storage media to another such as from the magnetic disc media 222 to the magnetic tape unit 224. Of Course, additional types and more than one type of storage media could be incorporated into the storage system 200.

Fig. 3 is a block diagram of another exemplary storage system 300 that is illustrated in greater detail than the storage system 100. The storage system 300 includes a computing device 302, a computing device 304, and a computing device 306 that are interconnected and communicate with one another on a network 307, such as an ethernet network. The computing device 306 includes

magnetic disc media 308 for storage of data that the computing device 306 receives from the computing device 302.

The computing device 302 includes software applications 310 and installed file system 312. When data is to be stored from the software application 310, the installed file system 312 begins interacting with a manager module 314 of the computing device 304. The manager module 314 includes storage policies 316, creation tools 318, and a master map 320. In this embodiment, when the software application is directed to stored data, the data is sent to the installed file system and then the manager module 314 examines the storage policies 316 to determine the appropriate location for storage of the data. The master map 320 includes further information for directing the data to be sent to the computing device 306 where a media module 322 receives the data and the data is then stored in the appropriate storage media. A data index 324 in the media module 322 contains still further information regarding the location to store the data. From computing device 306, the data may be stored in the magnetic disc media 308 or an optical media 326.

The master map 320 may direct the data to be stored in a computing device 328, rather than the computing device 306. In the event that data is to be stored in the computing device 328, a media module 330 determines exactly where the data is to be stored and updates a data index 332 when the data is stored in one of the storage media such as the optical media 326 or a magnetic tape media 334.

The master map 320 could also send the data to a computing device 336 where a media module 338 determines that the data will be placed at a different location and this information is updated in a data index 340 before the data is finally sent to other storage 342.

As illustrated in Fig. 3, the computing device 302 is used to store data from the software applications 310 in one of multiple storage locations. Of course, the storage media 308, 326, 334, and 342 could be located in other arrangements than as distributed between the computing devices 306, 328, 336. As those skilled in the art will understand upon viewing this disclosure, a manager module does not necessarily have to be located in the computing device 304, but, as illustrated in dashed lines, a manager module 344 could replace or supplement the manager module 314. Likewise, a manager module 346 illustrated in dashed lines could also replace or supplement the other manager modules 314 and 344. The media modules 322, 330, and 338 could also be supplemented with a media module 348 shown in dashed lines that would interact with an additional storage media 350, also illustrated in dashed lines. Further as illustrated in dashed lines, the computing device 304 could include a media module 352 that interacts with a storage media 354.

Fig. 4 is a block diagram of another exemplary storage system 400 that illustrates another embodiment of the storage system 100. The storage system 400 includes computing devices 402,

404, and 406, which communicate across network 407. Each of the computing devices 402, 404, and 406 communicate with a storage area network 408 and a network attached storage 410. The computing device 402 includes software applications 412 that communicate with an installed file system 414. The installed file system 414 communicates with a media module 416. In the event that the software applications 412 desire to store data, the data is transmitted through the installed file system 414 and to the media module 416. At the media module 416, the decision is made whether to send the data to the storage area network 408 or the network attached storage 410, or to some other location, such as storage media 418.

The computing device 404 illustrates a different configuration of a computing device and includes software applications 420 which communicate directly with an installed file system 422. When determining where the data is to be stored in the storage system 400, the installed file system 422 may communicate with the computing device 406. The computing device 406 includes a manager module 422 which includes storage policies 426, creation tools 428, and a master map 430. Based on the data that is to be stored, the installed file system 422 receives instructions from the manager module 424. The manager module 424 provides instructions that are at least based on the storage policies 426 and the master map 430. When the installed file system 422 has received instructions for where to store the data, the data is passed to the appropriate storage mechanism.

The storage area network 408 manages storage of data on a magnetic disk media 432, an optical media 434, and a magnetic tape media 436. Of course, storage area network 408 is an exemplary storage area network and could manage storage of data on numerous other types of storage media. The storage area network 408 communicates with the computing devices 402 and 404 via a high speed fibre network 437.

When the data is to be stored at the network attached storage 410, the installed file system 422 sends the data to the network attached storage 410 where a media module 438 directs the storage of the data at storage media 440. The storage media 440 represents one or more potential storage media that may be accessible for storage of the data.

As illustrated, in dashed lines, transmission of the data may occur in multiple paths. For example, a media module 442 may be present in the storage area network 408 or a media module 444 may be included in the manager module 424. In the event that the manager module 424 includes the media module 444, the fibre network 437 could be expanded as illustrated in dashed lines 446. In addition, the network 407 could be expanded to communicate directly with the storage area network 408 as illustrated by dashed line 448.

Fig. 5 is a block diagram of exemplary storage policies 106. The storage policies 106 include storage sequences 501 and storage logic 502. The storage sequences 501 represent potential storage patterns which may be used to store particular types of data. The storage logic 502 includes initial storage sequence selection 504, which represents potential storage considerations for the data when it arrives at the storage control 104. Represented in dash lines is storage sequence reselection 506 and storage sequence adaptation 508. The storage sequence reselection 506 represents logic which would allow alternative sequence selection logic to be used in place of the initial storage sequence selection 504. Although not required, a storage sequence reselection 506 would occur when the initial storage sequence selection 504 is inappropriate for the data that is being transmitted from the software application 102. The storage sequence adaptation 508 represents logic that would enable the initial storage sequence selection 504 to be modified according to changes in the data that is being saved in the storage mechanisms.

Fig. 6 is a block diagram of exemplary storage sequences 501. The storage sequences 501 include a primary storage sequence 602, a secondary storage sequence 604, a duplicate storage sequence 606, and a long term storage sequence 608. It should be noted that these are exemplary storage sequences only, and numerous other storage sequences could be added to the storage sequences 501. The storage sequences 501 are arranged in rows with columns to illustrate the different storage sequence characteristics. The far left column illustrates the storage sequence name followed by a first storage I.D. in the second column. For example, if the applications 102 are directing the storage of data, the storage control 104 would look to the storage policies 106. A default storage policy may include storing the data to the primary storage sequence 602 where the data would enter the first storage having an I.D. of 001. The data would be stored at the first storage I.D. for 13 weeks as indicated in the next column of the storage sequences 501. At this point the data would be moved to a second storage I.D. 005 where it would be stored for a duration of 26 weeks. After 26 weeks, the data would be moved to a third storage I.D. 002 for a duration of 52 weeks. This process would continue until the data is stored in an Nth storage I.D. 004. The primary storage sequence 602, of course, is an exemplary storage sequence, but is explained here as a means to understanding operations of the storage sequences 501. As can be seen from Fig. 6, the other storage sequences operate in a similar manner. Of particular note is the duplicate storage sequence 606.

In the duplicate storage sequence 606, data is received at the storage control 104 where it is determined that the data should be copied to two separate storage locations, which are represented in the first storage I.D. column as 001 and 005. The 001 storage I.D. has been directed to store the data for 13 weeks, and then the data may be deleted. The data in the storage having I.D. 005 is directed to

be stored for 39 weeks where it will then be moved to a second storage I.D. of 002 for a duration of 26 weeks, after which it will be moved to the nth storage I.D. 004.

Fig. 7 is a block diagram of the initial storage sequence selection 504. The initial storage sequence selection 504 is comprised of a number of elements that may be used to determine which of the storage sequences 501 that will be used to store the current data. A default initial storage sequence selection is selected for the storage system 100, and the default is comprised of a particular arrangement of factors such as a user directed override 702, a user profile 704, an application 706, a file type 708, a user network location 710, an available storage space 712, . . . . The user directed override 702 represents the situation where a user chooses to implement a particular storage sequence rather than the default storage sequence that would be used if the default storage policies were used. The user profile 704 represents logic that would decide which storage sequence to use based on the type of user, e.g., if the user is a manager, a certain storage sequence would be used, whereas a secretary may require a modified storage sequence. The application 706 that the data is generated from could also play a factor in which storage sequence to use, e.g., if the data was generated from a word processing application, one storage sequence could be selected, whereas if a video application generated the data, a different storage sequence could be selected. The file type 708 is commonly the most important factor in determining which storage sequence to use for the data that is received at the storage control 104. The user network location 710 could play a factor in which storage sequence to use, e.g., if the user is located near a particular storage device, one storage sequence may be desired over another storage sequence. The available storage space 712 could affect where data is to be stored if the default storage sequence requires that data be stored in a storage space that is reaching capacity and another storage space would be desirable because of extra space.

Fig. 8 is a block diagram of the storage sequence reselection 506. The storage sequence reselection 506 represents the factors that would be pertinent to selecting a new storage sequence for storage of data in the storage system 100. For example, if specific file usage history 802 shows that the data should be stored in a different storage sequence, the storage sequence reselection 506 allows the new selection of a storage sequence. The file type usage history 804, e.g., the access history of word documents, is another factor which may require a new storage sequence to be used in the storage system 100. Changes in user profile 806 are another reason for changing the storage sequences of the default system. Another factor for changing the default storage sequences is user network relocation 808. Still other factors include the available storage space 810 and added storage media 812. Of course, these are exemplary and additional factors could be added to the storage sequence reselection 506.

Fig. 9 is a block diagram of the storage sequence adaptation 508. Like the storage sequence reselection 506, the storage sequence adaptation 508 is illustrated having exemplary reasons for adapting the storage sequences that in the default storage control 104. A specific file usage history 902 could be a reason for changing a storage sequence, such as the primary storage sequence 602 having a greater duration for the first storage I.D. Changes in user profile 904 can also be a reason to adapt the storage sequences. Other factors for changing the storage sequences include user network relocation 906, available storage space 908, and added storage media 910.

**CLAIMS**

1. A computer storage system comprising:

a processor that supports operation of at least one software application that may be utilized to generate data that is stored in the computer storage system;

at least one storage media for storing selected data from the data that is generated;

a plurality of storage policies, each one of the plurality of storage policies including directions that the processor must follow for storage of data on the at least one storage media; and

a storage control that interacts with the at least one software application of the processor to select one of the plurality of storage policies and to determine, according to the selected one of the plurality of storage policies, which of the at least one storage media to use for storage of the selected data.

2. The computer storage system of claim 1 wherein the at least one storage media comprises a plurality of storage media that are selected from the group consisting of magnetic tape media, magnetic disk media, and optical medial.

3. The computer storage system of claim 1 wherein the storage control comprises a manager module that interacts with a media module to compile an index to track the location of the selected data in the at least one storage media.

4. The computer storage system of claim 1 wherein the storage control is configured to migrate the selected data among the at least one storage media according to the selected one of the plurality of storage policies.

5.. The computer storage system of claim 1 wherein the plurality of storage policies are selected from the group consisting of storage sequences, storage logic, initial storage sequence selection criteria, storage sequence reselection criteria, and storage sequence adaptation criteria.

6. The computer storage system of claim 5 wherein the initial storage sequence selection criteria are selected from the group consisting of user directed override, user profile, application, file type, user network location, and available storage space

7. The computer storage system of claim 5 wherein the storage sequence reselection criteria are selected from the group consisting of specific file usage history, file type usage history, user profile, user network relocation, available storage space, and added storage media.

8. The computer storage system of claim 5 wherein the storage sequence adaptation criteria are selected from the group consisting of specific file usage history, user profile, user network relocation, available storage space, and added storage media.

9. A computer storage system comprising:

a processor that supports operation of at least one software application in order to store selected data in the computer storage system;

at least one storage media for storing the selected data;

a plurality of storage policies, each having particular storage guidelines, that are available to determine how data is to be stored in the computer storage system; and

a storage control that interacts with the at least one software application of the processor to determine which of the plurality of storage policies to use for storage of the selected data and that stores the selected data according to the selected storage policy.

10. The computer storage system of claim 9 wherein the storage control migrates the selected data from one of the at least one storage media to another of the at least one storage media according to the selected storage policy.

11. A method for storing data on one of a plurality of storage media of a computer storage system according to one of a plurality of storage policies that are defined by a user, the method comprising:

directing a software application to store selected data;

examining the selected data for particular characteristics;

selecting a particular storage policy that matches the particular characteristics of the selected data; and

storing the selected data on a storage media of the computer storage system according to the particular storage policy.

12. The method of claim 11 further comprising migrating the selected data among the plurality storage media according to the particular storage policy.

1/9




Fig. 1

2/9




Fig. 2

3/9




Fig. 3

4/9




Fig. 4

5/9




Fig. 5

6/9

Storage Sequences						
Storage Name	Storage ID	Duration (weeks)	Storage ID	Duration (weeks)	Storage ID	Nth Storage ID
Primary	001	13	005	26	002	52
Secondary	005	13	003	13	002	65
Duplicate	001	13				004
	005	39	002	26		004
					⋮	
Long Term	005	2	002	12		004

Fig. 6

7/9




Fig. 7

8/9




Fig. 8

9/9




Fig. 9




(19) World Intellectual Property Organization  
International Bureau(43) International Publication Date  
2 August 2001 (02.08.2001)

PCT

(10) International Publication Number  
**WO 01/55856 A2**

- (51) International Patent Classification<sup>7</sup>: **G06F 12/00**
- (21) International Application Number: **PCT/US01/03183**
- (22) International Filing Date: 31 January 2001 (31.01.2001)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
 60/179,344 31 January 2000 (31.01.2000) US  
 09/774,268 30 January 2001 (30.01.2001) US
- (71) Applicant: **COMMVAULT SYSTEMS, INC. [US/US];**  
2 Crescent Place, Oceanport, NJ 07757-0090 (US).
- (72) Inventors: **IGNATIUS, Paul; 25 Highland Drive, Jackson, NJ 08527 (US). THEISEN, Marjorie, H.; 3272 S. Newcombe Street, Lakewood, CO 80227 (US). OSHINSKY, David, A.; 22 Francis Road, East Brunswick, NJ**
- 08816 (US). KAVURI, Srinivas; 40 Maple Court, Highland Park, NJ 08904 (US).
- (74) Agent: **SCOTT, Russell, C.; Akin, Gump, Strauss, Hauer & Feld, LLP, 816 Congress Avenue, 19th Floor, Austin, TX 78701 (US).**
- (81) Designated State (national): CA.
- (84) Designated States (regional): European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR).
- Published:  
 — without international search report and to be republished upon receipt of that report
- For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: LOGICAL VIEW AND ACCESS TO PHYSICAL STORAGE IN MODULAR DATA AND STORAGE MANAGEMENT SYSTEM

**WO 01/55856 A2**

(57) Abstract: A computer storage system having a processor that supports operation of at least one software application in order to store selected data in the computer storage system, at least one storage media for storing the selected data, a plurality of storage policies, each having particular storage guidelines, that are available to determine how data is to be stored in the computer storage system, and a storage control that interacts with the at least one software application of the processor to determine which of the plurality of storage policies to use for storage of the selected data and that stores the selected data according to the selected storage policy. In one embodiment, the storage control of the computer storage system migrates the selected data from one of the at least one storage media to another of the at least one storage media according to the selected storage policy.

**LOGICAL VIEW AND ACCESS TO PHYSICAL STORAGE IN  
MODULAR DATA AND STORAGE MANAGEMENT SYSTEM**

**I. SPECIFICATION**

**CROSS-REFERENCE TO RELATED APPLICATION**

The present application claims the benefit of U.S. Provisional Application Serial Number 60/179,344, filed January 31, 2000, pending, which is hereby incorporated by reference herein in its entirety.

**BACKGROUND**

**1. Field of the Invention**

The present invention generally relates to data storage in computer systems and more particularly to the logical view and access to physical storage in a modular data and storage management system.

**2. Description of the Related Art**

Existing storage systems require an end user to know and define exactly where data needs to be written and deposited in the storage system. This type of storage system is very hard to use and inherently unscalable as the user's data grows since the user will not be able to keep track of every storage device in the system. More importantly, every time new storage device(s) are added to the environment, the user will have to redefine how and where archival or backup data needs to be stored.

Many other problems and disadvantages of the prior art will become apparent to one skilled in the art after comparing such prior art with the present invention as described herein.

**SUMMARY**

Various aspects of the present invention may be realized through a computer storage system that includes a processor that supports operation of at least one software application that may be utilized to generate data that is stored in the computer storage system. The computer storage system also includes at least one storage media for storing selected data from the data that is generated. A plurality of storage policies are included where each one of the plurality of storage policies includes directions that the processor must follow for storage of data on the at least one storage media. A storage control interacts with the at least one software application of the processor to select one of the

plurality of storage policies and to determine, according to the selected one of the plurality of storage policies, which of the at least one storage media to use for storage of the selected data.

In certain embodiments, the at least one storage media of the computer storage system includes a plurality of storage media, e.g., magnetic tape media, magnetic disk media, optical medial, or other suitable storage media familiar to those skilled in the art. The storage control of the computer storage system may include a manager module that interacts with a media module to compile an index to track the location of the selected data in the at least one storage media. The storage control is often configured to migrate the selected data among the at least one storage media according to the selected one of the plurality of storage policies. The plurality of storage policies may be storage sequences, storage logic, initial storage sequence selection criteria, storage sequence reselection criteria, storage sequence adaptation criteria, etc. The initial storage sequence selection criteria are commonly user directed override, user profile, application, file type, user network location, and available storage space or similar criteria. The storage sequence reselection criteria are criteria such as specific file usage history, file type usage history, user profile, user network relocation, available storage space, added storage media, etc. Finally, the storage sequence adaptation criteria are items such as specific file usage history, user profile, user network relocation, available storage space, added storage media, or other suitable criteria that becomes apparent to those of ordinary skill in the art and viewing the present disclosure.

Various aspects of the present invention may also be realized by a computer storage system having a processor that supports operation of at least one software application in order to store selected data in the computer storage system, at least one storage media for storing the selected data, a plurality of storage policies, each having particular storage guidelines, that are available to determine how data is to be stored in the computer storage system, and a storage control that interacts with the at least one software application of the processor to determine which of the plurality of storage policies to use for storage of the selected data and that stores the selected data according to the selected storage policy. In one embodiment, the storage control of the computer storage system migrates the selected data from one of the at least one storage media to another of the at least one storage media according to the selected storage policy.

Still other aspects of the present invention are realized through a method for storing data on one of a plurality of storage media of a computer storage system according to one of a plurality of storage policies that are defined by a user. The method involves, not necessarily in this order directing a software application to store selected data; examining the selected data for particular characteristics; selecting a particular storage policy that matches the particular characteristics of the

selected data; and storing the selected data on a storage media of the computer storage system according to the particular storage policy.

The method may also include migrating the selected data among the plurality storage media according to the particular storage policy.

Other aspects of the present invention will become apparent with further reference to the drawings and specification which follow.

**BRIEF DESCRIPTION OF THE DRAWINGS**

Fig. 1 is a block diagram of an exemplary embodiment of a storage system that operates according to principles of the present invention.

Fig. 2 is a block diagram of an exemplary storage system illustrated in greater detail than the storage system of Fig. 1.

Fig. 3 is a block diagram of another exemplary storage system that is illustrated in greater detail than the storage system of Fig. 1.

Fig. 4 is a block diagram of another exemplary storage system that illustrates another embodiment of the storage system of Fig. 1.

Fig. 5 is a block diagram of exemplary storage policies as illustrated in Fig. 1.

Fig. 6 is a block diagram of exemplary storage sequences in greater detail than as illustrated in Fig. 5.

Fig. 7 is a block diagram of the initial storage sequence selection illustrated in greater detail than in Fig. 5.

Fig. 8 is a block diagram of the storage sequence reselection illustrated in greater detail than in Fig. 5.

Fig. 9 is a block diagram of the storage sequence adaptation illustrated in greater detail than in Fig. 5.

### DETAILED DESCRIPTION OF THE DRAWINGS

An exemplary description of the principles according to the present invention may be found in a storage policy where logical data buckets are presented to the end user for directing their data storage operations such as backup and archival. Storage policy is a logical concept. Each storage policy comprises one or more copies. Each copy is a self-contained unit of information. A primary copy indicates the default destination of storage operations and is tied to a practical set of drives, e.g., physical drives. Drives are addressed independently of the library or media agent that they are attached to. Storage policy at the same time comprises of a multiple number of streams which indicate the number of active drives that the storage operation can run against, allowing parallelism in data transfer. A single copy backup may always be run which transfers data to the primary copy. The primary copy may then be replicated to the secondary copies or a multiple copy simultaneous backup may be performed.

Each copy could have its destination media type different to allow storage policies to hide media type and related problems in a clean abstract interface. It also takes care of keeping track of which copy resides where, what is size of each data, etc. In this manner, the user's data is protected as well as presented in an abstract way using storage policies. Each of the copies of storage policy could be governed with different aging rules for the application's data. For example, the primary copy could be targeted to magnetic disc and aged every 30 days and secondary copy could be targeted to tape and aged every year. Data read operation will always fetch it from the optimal solution for retrieval.

In one embodiment of the present invention, the user just has to define a "Storage Policy" which defines where, how, and the duration the data should be stored at a higher level of abstraction without having to have intimate knowledge of the underlying storage and technology. The management of the details of data storage is transparent to the user.

Fig. 1 is a block diagram of an exemplary embodiment of a storage system 100 that operates according to principles of the present invention. The storage system 100 includes software applications 102 that a user interacts with to store data to various storage media. The software applications 102 communicate with a storage control 104 that determines where to store data from the software applications 102 according to storage policies 106. The storage control 104 typically receives data that is to be stored from the software applications 102. Upon receipt of data from the software applications 102, the storage control 104 determines whether to store the data in a first storage media 108, a second storage media 110, or any number of other storage media down to an Nth

storage media 112. The storage control 104 controls which storage media the data will be stored to as well as where data that has been previously stored will be moved to.

For example, the storage control 104 may move data that has been stored in the first storage media 108 into the second storage media 110 based on certain storage policies 106. In addition, the storage control 104 could move data from the second storage media 110 to any other of the storage media including the Nth storage media 112. Also, the storage control 104 could move data from the first storage media 108 into the Nth storage media 112. Of course, the data could be moved in either direction, i.e., the storage control 104 is capable of moving data between storage media.

Fig. 2 is a block diagram of an exemplary storage system 200 illustrated in greater detail than the storage system 100. The storage system 200 includes a single computing device 202 to carry out operations of the storage system 200. The computing device 202 includes a software application 204 that interacts with an installed file system 206 when a user desires to transfer data from the software application 204 to the storage media. For example, when a user saves data from the software application 204, the user is only aware that the data is being saved to the installed file system 206. The installed file system 206 interacts with a manager module 208, which in turn interacts with a media module 210.

The manager module 208 includes storage policies 212. The storage policies 212 determine how data that the manager module 208 received from the installed file system 206 will be stored in the storage media. The storage policies 212 interact with creation tools 214 to create the appropriate storage scenario for the data that the manager module 208 has received from the installed file system 206. The creation tools 214 may be accessed directly in the manager module 208, or remotely, as illustrated by remote creation tools 216. The manager module 208 also includes a master map 218 to assist in initial storage decisions in the storage media.

The media module 210 includes a data index 220 that includes further details of where the data is to be stored in the storage system 200. The data index 220 includes details such as the location of magnetic disc media 222 and magnetic tape media 224. The data index 220 is updated with file location information when any data is moved from one storage media to another such as from the magnetic disc media 222 to the magnetic tape unit 224. Of Course, additional types and more than one type of storage media could be incorporated into the storage system 200.

Fig. 3 is a block diagram of another exemplary storage system 300 that is illustrated in greater detail than the storage system 100. The storage system 300 includes a computing device 302, a computing device 304, and a computing device 306 that are interconnected and communicate with one another on a network 307, such as an ethernet network. The computing device 306 includes

magnetic disc media 308 for storage of data that the computing device 306 receives from the computing device 302.

The computing device 302 includes software applications 310 and installed file system 312. When data is to be stored from the software application 310, the installed file system 312 begins interacting with a manager module 314 of the computing device 304. The manager module 314 includes storage policies 316, creation tools 318, and a master map 320. In this embodiment, when the software application is directed to stored data, the data is sent to the installed file system and then the manager module 314 examines the storage policies 316 to determine the appropriate location for storage of the data. The master map 320 includes further information for directing the data to be sent to the computing device 306 where a media module 322 receives the data and the data is then stored in the appropriate storage media. A data index 324 in the media module 322 contains still further information regarding the location to store the data. From computing device 306, the data may be stored in the magnetic disc media 308 or an optical media 326.

The master map 320 may direct the data to be stored in a computing device 328, rather than the computing device 306. In the event that data is to be stored in the computing device 328, a media module 330 determines exactly where the data is to be stored and updates a data index 332 when the data is stored in one of the storage media such as the optical media 326 or a magnetic tape media 334.

The master map 320 could also send the data to a computing device 336 where a media module 338 determines that the data will be placed at a different location and this information is updated in a data index 340 before the data is finally sent to other storage 342.

As illustrated in Fig. 3, the computing device 302 is used to store data from the software applications 310 in one of multiple storage locations. Of course, the storage media 308, 326, 334, and 342 could be located in other arrangements than as distributed between the computing devices 306, 328, 336. As those skilled in the art will understand upon viewing this disclosure, a manager module does not necessarily have to be located in the computing device 304, but, as illustrated in dashed lines, a manager module 344 could replace or supplement the manager module 314. Likewise, a manager module 346 illustrated in dashed lines could also replace or supplement the other manager modules 314 and 344. The media modules 322, 330, and 338 could also be supplemented with a media module 348 shown in dashed lines that would interact with an additional storage media 350, also illustrated in dashed lines. Further as illustrated in dashed lines, the computing device 304 could include a media module 352 that interacts with a storage media 354.

Fig. 4 is a block diagram of another exemplary storage system 400 that illustrates another embodiment of the storage system 100. The storage system 400 includes computing devices 402,

404, and 406, which communicate across network 407. Each of the computing devices 402, 404, and 406 communicate with a storage area network 408 and a network attached storage 410. The computing device 402 includes software applications 412 that communicate with an installed file system 414. The installed file system 414 communicates with a media module 416. In the event that the software applications 412 desire to store data, the data is transmitted through the installed file system 414 and to the media module 416. At the media module 416, the decision is made whether to send the data to the storage area network 408 or the network attached storage 410, or to some other location, such as storage media 418.

The computing device 404 illustrates a different configuration of a computing device and includes software applications 420 which communicate directly with an installed file system 422. When determining where the data is to be stored in the storage system 400, the installed file system 422 may communicate with the computing device 406. The computing device 406 includes a manager module 422 which includes storage policies 426, creation tools 428, and a master map 430. Based on the data that is to be stored, the installed file system 422 receives instructions from the manager module 424. The manager module 424 provides instructions that are at least based on the storage policies 426 and the master map 430. When the installed file system 422 has received instructions for where to store the data, the data is passed to the appropriate storage mechanism.

The storage area network 408 manages storage of data on a magnetic disk media 432, an optical media 434, and a magnetic tape media 436. Of course, storage area network 408 is an exemplary storage area network and could manage storage of data on numerous other types of storage media. The storage area network 408 communicates with the computing devices 402 and 404 via a high speed fibre network 437.

When the data is to be stored at the network attached storage 410, the installed file system 422 sends the data to the network attached storage 410 where a media module 438 directs the storage of the data at storage media 440. The storage media 440 represents one or more potential storage media that may be accessible for storage of the data.

As illustrated, in dashed lines, transmission of the data may occur in multiple paths. For example, a media module 442 may be present in the storage area network 408 or a media module 444 may be included in the manager module 424. In the event that the manager module 424 includes the media module 444, the fibre network 437 could be expanded as illustrated in dashed lines 446. In addition, the network 407 could be expanded to communicate directly with the storage area network 408 as illustrated by dashed line 448.

Fig. 5 is a block diagram of exemplary storage policies 106. The storage policies 106 include storage sequences 501 and storage logic 502. The storage sequences 501 represent potential storage patterns which may be used to store particular types of data. The storage logic 502 includes initial storage sequence selection 504, which represents potential storage considerations for the data when it arrives at the storage control 104. Represented in dash lines is storage sequence reselection 506 and storage sequence adaptation 508. The storage sequence reselection 506 represents logic which would allow alternative sequence selection logic to be used in place of the initial storage sequence selection 504. Although not required, a storage sequence reselection 506 would occur when the initial storage sequence selection 504 is inappropriate for the data that is being transmitted from the software application 102. The storage sequence adaptation 508 represents logic that would enable the initial storage sequence selection 504 to be modified according to changes in the data that is being saved in the storage mechanisms.

Fig. 6 is a block diagram of exemplary storage sequences 501. The storage sequences 501 include a primary storage sequence 602, a secondary storage sequence 604, a duplicate storage sequence 606, and a long term storage sequence 608. It should be noted that these are exemplary storage sequences only, and numerous other storage sequences could be added to the storage sequences 501. The storage sequences 501 are arranged in rows with columns to illustrate the different storage sequence characteristics. The far left column illustrates the storage sequence name followed by a first storage I.D. in the second column. For example, if the applications 102 are directing the storage of data, the storage control 104 would look to the storage policies 106. A default storage policy may include storing the data to the primary storage sequence 602 where the data would enter the first storage having an I.D. of 001. The data would be stored at the first storage I.D. for 13 weeks as indicated in the next column of the storage sequences 501. At this point the data would be moved to a second storage I.D. 005 where it would be stored for a duration of 26 weeks. After 26 weeks, the data would be moved to a third storage I.D. 002 for a duration of 52 weeks. This process would continue until the data is stored in an Nth storage I.D. 004. The primary storage sequence 602, of course, is an exemplary storage sequence, but is explained here as a means to understanding operations of the storage sequences 501. As can be seen from Fig. 6, the other storage sequences operate in a similar manner. Of particular note is the duplicate storage sequence 606.

In the duplicate storage sequence 606, data is received at the storage control 104 where it is determined that the data should be copied to two separate storage locations, which are represented in the first storage I.D. column as 001 and 005. The 001 storage I.D. has been directed to store the data for 13 weeks, and then the data may be deleted. The data in the storage having I.D. 005 is directed to

be stored for 39 weeks where it will then be moved to a second storage I.D. of 002 for a duration of 26 weeks, after which it will be moved to the nth storage I.D. 004.

Fig. 7 is a block diagram of the initial storage sequence selection 504. The initial storage sequence selection 504 is comprised of a number of elements that may be used to determine which of the storage sequences 501 that will be used to store the current data. A default initial storage sequence selection is selected for the storage system 100, and the default is comprised of a particular arrangement of factors such as a user directed override 702, a user profile 704, an application 706, a file type 708, a user network location 710, an available storage space 712, ... . The user directed override 702 represents the situation where a user chooses to implement a particular storage sequence rather than the default storage sequence that would be used if the default storage policies were used. The user profile 704 represents logic that would decide which storage sequence to use based on the type of user, e.g., if the user is a manager, a certain storage sequence would be used, whereas a secretary may require a modified storage sequence. The application 706 that the data is generated from could also play a factor in which storage sequence to use, e.g., if the data was generated from a word processing application, one storage sequence could be selected, whereas if a video application generated the data, a different storage sequence could be selected. The file type 708 is commonly the most important factor in determining which storage sequence to use for the data that is received at the storage control 104. The user network location 710 could play a factor in which storage sequence to use, e.g., if the user is located near a particular storage device, one storage sequence may be desired over another storage sequence. The available storage space 712 could affect where data is to be stored if the default storage sequence requires that data be stored in a storage space that is reaching capacity and another storage space would be desirable because of extra space.

Fig. 8 is a block diagram of the storage sequence reselection 506. The storage sequence reselection 506 represents the factors that would be pertinent to selecting a new storage sequence for storage of data in the storage system 100. For example, if specific file usage history 802 shows that the data should be stored in a different storage sequence, the storage sequence reselection 506 allows the new selection of a storage sequence. The file type usage history 804, e.g., the access history of word documents, is another factor which may require a new storage sequence to be used in the storage system 100. Changes in user profile 806 are another reason for changing the storage sequences of the default system. Another factor for changing the default storage sequences is user network relocation 808. Still other factors include the available storage space 810 and added storage media 812. Of course, these are exemplary and additional factors could be added to the storage sequence reselection 506.

Fig. 9 is a block diagram of the storage sequence adaptation 508. Like the storage sequence reselection 506, the storage sequence adaptation 508 is illustrated having exemplary reasons for adapting the storage sequences that in the default storage control 104. A specific file usage history 902 could be a reason for changing a storage sequence, such as the primary storage sequence 602 having a greater duration for the first storage I.D. Changes in user profile 904 can also be a reason to adapt the storage sequences. Other factors for changing the storage sequences include user network relocation 906, available storage space 908, and added storage media 910.

**CLAIMS**

1. A computer storage system comprising:

a processor that supports operation of at least one software application that may be utilized to generate data that is stored in the computer storage system;

at least one storage media for storing selected data from the data that is generated;

a plurality of storage policies, each one of the plurality of storage policies including directions that the processor must follow for storage of data on the at least one storage media; and

a storage control that interacts with the at least one software application of the processor to select one of the plurality of storage policies and to determine, according to the selected one of the plurality of storage policies, which of the at least one storage media to use for storage of the selected data.

2. The computer storage system of claim 1 wherein the at least one storage media comprises a plurality of storage media that are selected from the group consisting of magnetic tape media, magnetic disk media, and optical medial.

3. The computer storage system of claim 1 wherein the storage control comprises a manager module that interacts with a media module to compile an index to track the location of the selected data in the at least one storage media.

4. The computer storage system of claim 1 wherein the storage control is configured to migrate the selected data among the at least one storage media according to the selected one of the plurality of storage policies.

5. The computer storage system of claim 1 wherein the plurality of storage policies are selected from the group consisting of storage sequences, storage logic, initial storage sequence selection criteria, storage sequence reselection criteria, and storage sequence adaptation criteria.

6. The computer storage system of claim 5 wherein the initial storage sequence selection criteria are selected from the group consisting of user directed override, user profile, application, file type, user network location, and available storage space

7. The computer storage system of claim 5 wherein the storage sequence reselection criteria are selected from the group consisting of specific file usage history, file type usage history, user profile, user network relocation, available storage space, and added storage media.

8. The computer storage system of claim 5 wherein the storage sequence adaptation criteria are selected from the group consisting of specific file usage history, user profile, user network relocation, available storage space, and added storage media.

9. A computer storage system comprising:

a processor that supports operation of at least one software application in order to store selected data in the computer storage system;

at least one storage media for storing the selected data;

a plurality of storage policies, each having particular storage guidelines, that are available to determine how data is to be stored in the computer storage system; and

a storage control that interacts with the at least one software application of the processor to determine which of the plurality of storage policies to use for storage of the selected data and that stores the selected data according to the selected storage policy.

10. The computer storage system of claim 9 wherein the storage control migrates the selected data from one of the at least one storage media to another of the at least one storage media according to the selected storage policy.

11. A method for storing data on one of a plurality of storage media of a computer storage system according to one of a plurality of storage policies that are defined by a user, the method comprising:

directing a software application to store selected data;

examining the selected data for particular characteristics;

selecting a particular storage policy that matches the particular characteristics of the selected data; and

storing the selected data on a storage media of the computer storage system according to the particular storage policy.

12. The method of claim 11 further comprising migrating the selected data among the plurality storage media according to the particular storage policy.




Fig. 1




Fig. 2




Fig. 3




Fig. 4




Fig. 5

*s01*

Storage Sequences						
Storage Name	1st Storage ID	Duration (weeks)	2nd Storage ID	Duration (weeks)	3rd Storage ID	Duration (weeks)
Primary	001	13	005	26	002	52
Secondary	005	13	003	13	002	65
Duplicate	001	13	005	39	002	26
Long Term	005	2	002	12	004	

*102*      *104*      *106*      *108*

• • •

Fig. 6




Fig. 7




Fig. 8




Fig. 9

